

1-12-00

A

Case Docket No. PHA 23,656

THE COMMISSIONER OF PATENTS AND TRADEMARKS, Washington, D.C. 20231

Enclosed for filing is the patent application of Inventor(s):

KARL M. ISHAM

For: METHOD FOR DYNAMIC LOANING IN RATE MONOTONIC REAL-TIME SYSTEMS

ENCLOSED ARE:

- ☒ Associate Power of Attorney;
☒ Information Disclosure Statement, Form PTO-1449 and copies of documents listed therein;
☐ Preliminary Amendment;
☒ Specification (20 Pages of Specification, Claims, & Abstract);
☒ Declaration and Power of Attorney:
 (2 Pages of a ☒ fully executed ☐ unsigned Declaration);
☒ Drawing (2 sheets of ☒ informal ☐ formal sheets);
☐ Certified copy of application Serial No. ;
☒ Other: Authorization Pursuant to 37 CFR 1.136
☒ Assignment to Philips Electronics North America Corporation

FEE COMPUTATION

CLAIMS AS FILED				
FOR	NUMBER FILED	NUMBER EXTRA	RATE	BASIC FEE - 690.00
Total Claims	25- 20 =	5	X \$18 =	90.00
Independent Claims	6- 3 =	3	X \$78 =	234.00
Multiple Dependent Claims, if any			\$260 =	0.00
TOTAL FILING FEE			=	\$1,014.00

Please charge Deposit Account No. 14-1270 in the amount of the total filing fee indicated above, plus any deficiencies. The Commissioner is also hereby authorized to charge any other fees which may be required, except the issue fee, or credit any overpayment to Account No. 14-1270.

☐ Amend the specification by inserting before the first line the sentence: --This is a continuation-in-part of application Serial No. , filed .--.

CERTIFICATE OF EXPRESS MAILINGExpress Mail Mailing Label No. EL335549502USDate of Deposit January 11, 2000

I hereby certify that this paper and/or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231.

Elissa DeLuccy
Printed Name

Elissa DeLuccy
Signature

[Signature]
Tony Piotrowski, Reg. No. 42,080
Attorney
(914) 333-9609
U.S. Philips Corporation
580 White Plains Road
Tarrytown, New York 10591
F:\WPDOCS\pi\mp21pim0.nc0.doc

01/11/00
JCS44 U.S. PTO

09434374 044100

09/481771
JCS675 U.S. PTO
01/11/00

METHOD FOR DYNAMIC LOANING IN RATE MONOTONIC REAL-TIME SYSTEMS**Cross Reference to Related Applications**

5 This application claims the benefit of United States
Provisional Application Number 60/129,301, filed April 15, 1999.

Field of the Invention

10 The present invention relates to the timing behavior of
real-time computing systems, and more particularly, to a method
and apparatus for loaning execution capacity among tasks
executing in a real-time computing system.

Background of the Invention

15 Real-time systems differ from other forms of computing
in that they must be temporally correct as well as logically
correct. Such systems are developed to satisfy the following
three primary criteria: guaranteed timing deadlines, fast
response times, and stability in overload. Schedulability
20 describes the capacity of a system. A system that is schedulable
can meet all of its critical timing deadlines. Latency describes
the responsiveness of a system. In a real-time system, it is the
worst-case system response time to events that matters.
Stability in overload means the system is able to meet its
25 critical deadlines even if all deadlines cannot be met.

 One of the most useful models available for developing
real-time computing systems is Rate Monotonic Analysis (RMA).
RMA provides a mathematical framework for reasoning about system
timing behavior and provides an engineering basis for designing
30 real-time systems. RMA was first disclosed in Liu & Layland,
"Scheduling Algorithms for Multi-Programming in a Hard Real-Time
Environment," Journal of the Ass'n of Computing Machinery (ACM)
20, 1, 40-61 (January, 1973), incorporated by reference herein.
Generally, Liu and Layland demonstrated that a set of n periodic

tasks with deadlines at the end of their periods will meet their deadlines if they are arranged in priority according to their periods, and they meet a schedulability bound test. Since this paper, RMA has evolved into a collection of methods that have
5 extended the original theory from its original form. The basic RMA real-time guarantee, however, has remained unchanged. For a general discussion of these collections of methods, see, for example, Klein et al, A Practitioner's Handbook for Real-Time Analysis: Guide to Rate Monotonic Analysis for Real-time Systems
10 (Kluwer Academic Publishing, ISBN 0-7923-9361-9, 1993), incorporated by reference herein.

Methods currently exist to assess spare capacity and dynamically change the priority of tasks to alter the scheduling policy of the system. Spare capacity is the amount of execution
15 time that can be added to the response of an event while preserving the schedulability of lower priority events. A related method, eliminating overrun, computes the amount of resource usage that must be eliminated to allow an event to meet its deadlines. See, Klein et al, A Practitioner's Handbook for Real-Time Analysis: Guide to Rate Monotonic Analysis for Real-time Systems, Chapter 4, Group 3 (Kluwer Academic Publishing,
20 ISBN 0-7923-9361-9, 1993). Changing the priority of a task is used in various synchronization protocols to avoid priority inversion when multiple tasks share common data.

25 The validity of Rate Monotonic Analysis depends on preparing for the worst-case in terms of individual event timing and the concurrence of events. In other words, the system will be capable of meeting all of its deadlines if all worst-case events can be handled simultaneously. While mathematically
30 certain, assuming worst-case timings and coincidences required by the RMA analysis may result in harsh feature/performance tradeoffs for the usually limited hardware computing capacity found in a typical real-time system. While the worst-case high

priority event may occur very infrequently, the execution time for the worst-case high priority event must still be accounted for in RMA calculations. The capacity allocated to handle this rare worst-case event at high priority then gets traded off against the capacity available to lower priority events.

If RMA calculations show that the system is no longer schedulable with these maximum execution requirements, the designer is left with few options. Usually, the worst-case events will have to be redesigned to reduce their execution times. If the execution times cannot be reduced, the designer is left with discounting the offending event as an overload, provided the occasional timing violations to lower priority events can be tolerated. An overload condition is an acceptable alternative only if all lower priority events have soft deadlines (can tolerate a missed deadline occasionally). Unfortunately, this is rarely the case. In a system where such an overload situation exists, the real-time criteria of stability in overload is violated.

Currently, RMA only provides the tools necessary to determine if a system is schedulable. In other words, RMA will only tell you if a given design will work or not. For example, the method of assessing spare capacity is useful for determining how much execution time can be safely added to an event before timing requirements are broken. Likewise, RMA provides a method for determining how much execution time must be eliminated to meet timing requirements. These methods will prove useful for providing target values for redesign, but will not help if the execution times are intractable. There is no current method that allows a designer to gracefully handle intractable overloads.

Summary of the Invention

Generally, a method and apparatus are disclosed for sharing execution capacity among tasks executing in a real-time

computing system by pairing a high priority task having hard deadlines with a lower priority task having soft deadlines. The present invention extends RMA techniques for characterizing system timing behavior and designing real-time systems. During an overload condition, the higher priority task can dynamically borrow execution time from the execution capacity of the lower priority task without affecting the schedulability of the rest of the system. The higher priority task is bolstered in proportion to the capacity borrowed from the lower priority task, so that the combined utilization of the two tasks remains constant.

According to another aspect of the invention, the period of the degraded task is increased to compensate for the execution time that was loaned to the higher priority task. Thus, events for the degraded task may still be assigned according to the original execution budget, but are allowed a longer time to complete due to the increased work in the borrowing higher priority task. In addition, the priority of the lower priority task is modified to match the new period, in accordance with the Rate Monotonic Scheduling algorithm.

The present invention isolates effects of an overload to a particular task (i.e., the degraded task). In addition, the manner in which execution capacity is borrowed does not diminish the original execution budget of the degraded task, since the period of the degraded task is lengthened to compensate for the borrowed time. Thus, the overloaded system takes longer to perform non-critical events but the amount of work it will accept at any one time remains the same. In addition, the present invention provides better utilization of computing resources by borrowing non-overload execution time against the capacity of a rarely used task. In this case, a low priority task may be prepared to handle an intermittent event with soft deadlines. The present invention permits a higher priority task to borrow

some of the capacity for its normal operation from the intermittent task.

A more complete understanding of the present invention, as well as further features and advantages of the present invention, will be obtained by reference to the following detailed description and drawings.

Brief Description of the Drawings

FIG. 1 illustrates a real-time computing system in accordance with the present invention; and

FIG. 2 illustrates a state diagram of the capacity loaning process of FIG. 1.

Detailed Description

The present invention provides a method for pairing a high priority task having hard deadlines with a lower priority task with soft deadlines. During an overload, the higher priority task can dynamically borrow execution time from the execution capacity of the lower priority task without affecting the schedulability of the rest of the system. In this manner, the higher priority task can be bolstered in a proportional manner, so that the combined utilization of the two tasks remains constant.

According to another aspect of the present invention, the period of the degraded loaning task is lengthened to compensate for the borrowed time. In addition, the priority of the degraded task is modified to match the new period. Furthermore, the performance of the loaning task is degraded for the period of the borrowing task, extended over the whole period of the loaning task.

The present invention provides a means to gracefully degrade system performance during overload while still providing critical schedulability guarantees. For intractable overloads,

this may be the only practical method of maintaining stability. Secondly, the task to be degraded can be identified during the design stage, which is an important consideration from the standpoint of predictability and isolation. Real-time problems are notorious for their ability to elude capture in the field. If an overload event coincided with a different lower priority event every time it manifested itself, it would be nearly impossible to correctly identify and correct.

The present invention isolates the effects of an overload to a particular task. In addition, the means used by the present invention to borrow execution capacity does not diminish the original execution budget of the degraded task, since the period of the degraded task is lengthened to compensate for the borrowed time. Thus, the overloaded system takes longer to perform non-critical events but the amount of work it will accept at any one time remains the same. Finally, the method of the present invention can be used to provide better utilization of computing resources by borrowing non-overload execution time against the capacity of a rarely used task. In this case, a low priority task may be prepared to handle an intermittent event with soft deadlines. The present invention permits a higher priority task to borrow some of the capacity for its normal operation from the intermittent task.

RMA PRINCIPLES AND TERMINOLOGY

Schedulability Bound Test

As previously indicated, Liu and Layland demonstrated that a set of n periodic tasks with deadlines at the end of their periods will meet their deadlines if they are arranged in priority according to their periods, and they meet a schedulability bound test. For a detailed discussion of the schedulability bound test, see Liu & Layland, "Scheduling Algorithms for Multi-Programming in a Hard Real-Time

Environment," Journal of the Ass'n of Computing Machinery (ACM) 20, 1, 40-61 (January, 1973), incorporated by reference herein.

Generally, the schedulability bound test states that a set of n independent periodic tasks scheduled by the rate monotonic algorithm will always meet its deadlines, for all task phasings, if:

$$\frac{C_1}{T_1} + \dots + \frac{C_n}{T_n} \leq U(n) = n \left(2^{\frac{1}{n}} - 1 \right)$$

where,

C_i = worst-case task execution time of task $_i$,

T_i = period of task $_i$, and

$U(n)$ = utilization for n tasks.

Mathematical Basis for this Invention

As previously indicated, the present invention pairs two tasks, a critical task with hard deadlines and another task of lower priority with soft deadlines. Consider a pair of tasks whose total utilization remains constant, thereby allowing RMA to be applied to guarantee schedulability. Such a pair might exist in a real-time system in order to satisfy two different quality of service levels. The task pair relationship can be expressed as follows:

$$\frac{C_u}{T_u} + \frac{C_r}{T_r} = U \quad \text{Equation (1)}$$

where,

C_u = worst-case task execution time of task $_u$,

T_u = period of task $_u$,

C_r = worst-case task execution time of task $_r$,

T_r = period of task $_r$, and

U = utilization for both tasks.

Thus, the utilization of one task in the pair may be increased if the utilization of the other task in the pair is proportionally decreased to maintain a constant utilization, U .

In this manner, execution time can be borrowed from one task to supplement the execution of the other task. The techniques of the present invention are useful, for example, where events are validated against the available execution time, C_u or C_r , before they are assigned to a servicing task, $task_u$ or $task_r$. It follows that temporary overloads may be assigned to the higher priority task without sacrificing RMA guarantees. The utilization loan can be expressed by the equation:

$$\frac{C_u + N_u}{T_u} + \frac{C_r - N_r}{T_r} = U$$

$$\frac{C_u}{T_u} + \frac{N_u}{T_u} + \frac{C_r}{T_r} - \frac{N_r}{T_r} = U$$

$$\frac{N_u}{T_u} - \frac{N_r}{T_r} = U - \frac{C_u}{T_u} - \frac{C_r}{T_r}$$

$$\frac{N_u}{T_u} - \frac{N_r}{T_r} = 0 \quad (\text{from equation (1)})$$

$$\frac{N_u}{T_u} = \frac{N_r}{T_r}$$

$$N_u = \frac{N_r T_u}{T_r} \quad \text{Equation (2)}$$

where,

N_r = amount of execution time to borrow from $task_r$,
where $N_r < C_r$; and

N_u = amount of execution time available to loan to $task_u$.

The utilization terms for the two tasks in question are isolated from the utilization bound equation and equated to their combined utilization. In the above discussion, $task_u$ refers to the higher priority (urgent) task while $task_r$ refers to the lower priority (routine) task with soft deadlines. For the RMA utilization bound test to remain valid, the sum of these two tasks must remain constant during the application of the method.

Given that the utilization, U , remains constant, the utilization of the higher priority task may be increased as long

as there is a proportional decrease in the utilization of the lower priority task. The proportional change guarantees the schedulability of the rest of the system, since the net effect on the utilization bound equation is the same before and after the change. Equation (2) applies this proportional change and is solved for the amount of execution time the higher priority task may borrow.

While adding execution time to a task may have an application in handling overloads, it is usually not practical to actually remove execution time from another task to compensate for it. The execution times of most events at runtime are fixed, and therefore cannot be limited. This is solved by assuming that the original execution budget is fixed and letting the period of the task change instead. As discussed hereinafter, Equations (3) and (4) deal with this change in the loaning task's period. Since the period of a task also determines its priority in a Rate Monotonic environment, the priority of the task must change as well, using the usual set priority functions available in most real-time operating systems.

To compensate for the borrowed execution time, the period of the loaning task may be changed instead of limiting the execution time of the task. In this manner, the level of service in the loaning task is degraded gracefully. Events may still be assigned according to the original execution budget, C_r , but are allowed a longer time to complete due to the increased work in the borrowing task. The graceful degradation may be implemented as follows:

$$\frac{C_r}{T_n} = \frac{C_r - N_r}{T_r}$$

$$T_n = \frac{C_r}{\left(\frac{C_r - N_r}{T_r} \right)}$$

$$T_n = \frac{Cr \cdot Tr}{Cr - Nr} \quad \text{Equation (3)}$$

where

T_n = the new period of task_r.

To prevent an unbounded rise in the execution period of the task, T_n , Nr must be limited to a maximum loan amount where $Nr \ll Cr$. If it is assumed that at its maximum value, T_n , is a multiple of Tr , then:

$$m \cdot Tr = \frac{Cr \cdot Tr}{Cr - Nm}$$

$$m = \frac{Cr \cdot Tr}{(Cr - Nm) \cdot Tr}$$

$$m = \frac{Cr}{Cr - Nm}$$

$$Cr - Nm = \frac{Cr}{m}$$

$$Nm = Cr - \frac{Cr}{m}$$

$$Nm = Cr \left(1 - \frac{1}{m} \right) \quad \text{Equation (4)}$$

where

m = multiple of the period of task_r, and

Nm = maximum execution time, Nr , loanable from task_r.

FIG. 1 illustrates a real-time computing system 100 in accordance with the present invention. As shown in FIG. 1, the real-time computing system 100 includes certain standard hardware components, such as a processor 110 and a data storage device 120, such as a read-only memory and/or a random access memory (RAM).

The data storage device 120 includes a capacity loaning process 200, discussed further below in conjunction with FIG. 2.

Generally, the capacity loaning process 200 implements capacity loaning between two tasks in accordance with the present

invention. As shown in FIG. 1, the data storage device 120 also includes an operating system 150 that, among other things, manages the pairing of the two tasks, $task_U$ and $task_R$, for capacity loaning in accordance with the present invention.

FIG. 2 illustrates a state diagram of the capacity loaning process 200. As shown in FIG. 2, each task is modeled here as a sporadic server, with modifications due to the capacity loan algorithm shown in bold type. For a discussion of the UML notations used herein, see, for example, UML Notation Guide, version 1.1 (Sept. 1, 1997), downloadable from <http://www.omg.org/docs/ad/97-08-05.pdf>, and incorporated by reference herein.

The capacity loaning process 200 utilizes two assumptions. First, the amount of borrowed time required by the urgent task will be known at the beginning of the execution period of the task. This implies that each event must be marked with its worst-case execution time so a dynamic assessment of budget can be made by the task at runtime. Secondly, the loan is only viable during the period of the borrower.

At system initialization, both tasks must agree on the maximum capacity that will be made available to loan. Using Equation (4), the routine task would compute the maximum amount of execution time to loan based on its original execution budget and a limiting factor to prevent an unbounded rise in its period. This loanable amount would then be registered with the urgent task, which would then compute the maximum amount of execution time it may borrow with Equation (2) (accounting for the differences in the loan amount due to the differing period of the two tasks).

At the beginning of each execution period, the urgent task computes the amount of execution time over budget during state 220 and sends a message 240 to the routine task containing this amount (which may be 0, in the event that it is not over

budget). The urgent task then executes these events normally during state 220.

Because the loan is only viable during the period of the borrower, the routine task maintains two state variables.

5 The amount of capacity requested is the amount of execution time that has yet to be acted on. That is, this borrowed time has not yet degraded the performance of the loaner. The amount of capacity on loan is borrowed time that is currently acting to degrade the loaner. The action of the routine task in response to
10 the message 240 from the urgent task depends on the task's current state. While waiting during state 250, the overbudget amount from the urgent task is applied to capacity requested. Otherwise, the overbudget amount is added to capacity on loan, where it is then used with Equation (3) to set a new period for
15 the task. The priority of the task must also be lowered to match the new task period. At the beginning of the execution period of the routine task during state 260, capacity on loan is set from capacity requested, and Equation (3) is used to compute the task's period. Whenever the period of the task is changed, the
20 corresponding priority must be changed accordingly, to satisfy the Rate Monotonic scheduling algorithm.

It is to be understood that the embodiments and variations shown and described herein are merely illustrative of the principles of this invention and that various modifications
25 may be implemented by those skilled in the art without departing from the scope and spirit of the invention.

Claims

What is claimed is:

- 5 1. A method for sharing execution capacity among tasks executing in a real-time computing system having a performance specification in accordance with Rate Monotonic Analysis (RMA), comprising the steps of:

10 pairing a higher priority task with a lower priority task;

reallocating execution time from the lower priority task to the higher priority task during an overload condition; and

15 increasing the period of the lower priority task to compensate for said reallocated execution time.

- 20 2. The method of claim 1, wherein an amount of said execution time available to loan from said lower priority task, $task_R$, to said higher priority task, $task_u$, is obtained as follows:

$$Nu = \frac{Nr \cdot Tu}{Tr}$$

where,

N_r = amount of execution time to borrow from $task_R$,
where $N_r < C_r$,

25 T_r = period of $task_R$, and
 T_u = period of $task_u$.

3. The method of claim 1, wherein said increased period of the lower priority task, $task_R$, is obtained as follows:

30
$$Tn = \frac{Cr \cdot Tr}{Cr - Nr}$$

where

C_r = worst-case task execution time of task_r,

T_r = period of task_r, and

N_r = amount of execution time to borrow from task_r,

where $N_r < C_r$.

5

4. The method of claim 1, further comprising the step of limiting an amount of execution time, N_r , to borrow from said lower priority task, task_r, to a maximum loan amount where $N_r \ll C_r$, where

10

C_r = worst-case task execution time of task_r, and

N_r = amount of execution time to borrow from task_r.

5. The method of claim 4, wherein a maximum execution time, N_m , that may be borrowed from said lower priority task, task_r, is obtained as follows:

15

$$N_m = C_r \left(1 - \frac{1}{m} \right)$$

where m is the multiple of the period of said lower priority task, task_r.

20

6. The method of claim 1, wherein said higher priority task has hard deadlines.

7. The method of claim 1, wherein said lower priority task has soft deadlines.

25

8. A method for allocating resources among tasks executing in a real-time computing system having a performance specification in accordance with Rate Monotonic Analysis (RMA), comprising the steps of:

30

pairing a higher priority task with a lower priority task;

providing a first resource allocation to said lower priority task during a normal operating condition; and

reallocating a portion of said first resource allocation from said lower priority task to said higher priority task when said higher priority task is operable.

9. The method of claim 8, wherein said reallocated portion of said first resource allocation is obtained as follows:

$$Nu = \frac{Nr \cdot Tu}{Tr}$$

10 where,

N_r = amount of execution time to borrow from task_r,
where $N_r < C_r$,

T_r = period of task_r, and

T_u = period of task_u.

10. The method of claim 8, further comprising the step of increasing a period of said lower priority task, task_r, as follows:

$$Tn = \frac{Cr \cdot Tr}{Cr - Nr}$$

20 where

C_r = worst-case task execution time of task_r,

T_r = period of task_r, and

N_r = amount of execution time to borrow from task_r,
where $N_r < C_r$.

11. The method of claim 8, further comprising the step of limiting an amount of execution time, N_r , to reallocate from said lower priority task, task_r, to a maximum loan amount where $N_r < C_r$, where

C_r = worst-case task execution time of task_r, and

N_r = amount of execution time to borrow from task_r.

12. The method of claim 11, wherein a maximum execution time, Nm , that may be borrowed from said lower priority task, $task_r$, is obtained as follows:

$$Nm = Cr \left(1 - \frac{1}{m} \right)$$

where m is the multiple of the period of said lower priority task, $task_r$.

13. The method of claim 8, wherein said higher priority task has hard deadlines.

14. The method of claim 8, wherein said lower priority task has soft deadlines.

15. A method for sharing execution capacity among tasks executing in a real-time computing system having a performance specification in accordance with Rate Monotonic Analysis (RMA), comprising the steps of:

pairing a higher priority task, $task_u$, with a lower priority task, $task_r$;

reallocating execution time from the lower priority task to the higher priority task during an overload condition; and

increasing the utilization of said higher priority task; and

decreasing the utilization of said lower priority task in a proportional manner to maintain a constant utilization, U .

16. The method of claim 15, wherein said utilizations of said tasks are varied as follows:

$$\frac{Cu}{Tu} + \frac{Cr}{Tr} = U$$

where,

C_u = worst-case task execution time of task_u,

T_u = period of task_u,

C_r = worst-case task execution time of task_r,

5 T_r = period of task_r, and

U = utilization for both tasks.

17. The method of claim 15, wherein an amount of said execution time available to reallocate from said lower priority task, task_r, to said higher priority task, task_u, is obtained as follows:

$$N_u = \frac{N_r T_u}{T_r}$$

where,

N_r = amount of execution time to borrow from task_r,

15 where $N_r < C_r$,

T_r = period of task_r, and

T_u = period of task_u.

18. The method of claim 15, further comprising the step of increasing a period of the lower priority task, task_r, as follows:

$$T_n = \frac{C_r \cdot T_r}{C_r - N_r}$$

where

C_r = worst-case task execution time of task_r,

25 T_r = period of task_r, and

N_r = amount of execution time to borrow from task_r,

where $N_r < C_r$.

19. The method of claim 15, further comprising the step of limiting an amount of execution time, N_r , to borrow from said

lower priority task, $task_r$, to a maximum loan amount where $N_r < C_r$, where

C_r = worst-case task execution time of $task_r$, and
 N_r = amount of execution time to borrow from $task_r$.

5

20. The method of claim 19, wherein a maximum execution time, N_m , that may be borrowed from said lower priority task, $task_r$, is obtained as follows:

$$N_m = C_r \left(1 - \frac{1}{m} \right)$$

10 where m is the multiple of the period of said lower priority task, $task_r$.

21. The method of claim 15, wherein said higher priority task has hard deadlines.

22. The method of claim 15, wherein said lower priority task has soft deadlines.

23. A real-time computing system having a performance specification in accordance with Rate Monotonic Analysis (RMA), comprising:

a memory for storing computer readable code; and

a processor operatively coupled to said memory, said processor configured to:

25 pair a higher priority task with a lower priority task;
 reallocate execution time from the lower priority task to the higher priority task during an overload condition; and
 increase the period of the lower priority task to compensate for said reallocated execution time.

30

24. A real-time computing system having a performance specification in accordance with Rate Monotonic Analysis (RMA), comprising:

a memory for storing computer readable code; and

a processor operatively coupled to said memory, said processor configured to:

pair a higher priority task with a lower priority task;

provide a first resource allocation to said lower priority task during a normal operating condition; and

reallocate a portion of said first resource allocation from said lower priority task to said higher priority task when said higher priority task is operable.

25. A real-time computing system having a performance specification in accordance with Rate Monotonic Analysis (RMA), comprising:

a memory for storing computer readable code; and

a processor operatively coupled to said memory, said processor configured to:

pair a higher priority task, $task_u$, with a lower priority task, $task_r$;

reallocate execution time from the lower priority task to the higher priority task during an overload condition; and

increase the utilization of said higher priority task;

and
decrease the utilization of said lower priority task in a proportional manner to maintain a constant utilization, U .

ABSTRACT

5 A method and apparatus are disclosed for sharing
execution capacity among tasks executing in a real-time computing
system. The present invention extends RMA techniques for
characterizing system timing behavior and designing real-time
systems. A high priority task having hard deadlines is paired
with a lower priority task having soft deadlines. During an
overload condition, the higher priority task can dynamically
10 borrow execution time from the execution capacity of the lower
priority task without affecting the schedulability of the rest of
the system. The higher priority task is bolstered in a
proportion to the capacity borrowed from the lower priority task,
so that the combined utilization of the two tasks remains
15 constant. The period of the degraded task is increased to
compensate for the execution time that was loaned to the higher
priority task. In addition, the priority of the lower priority
task is modified to match the new period.

20 1100-19.app

REAL-TIME COMPUTING SYSTEM - 100

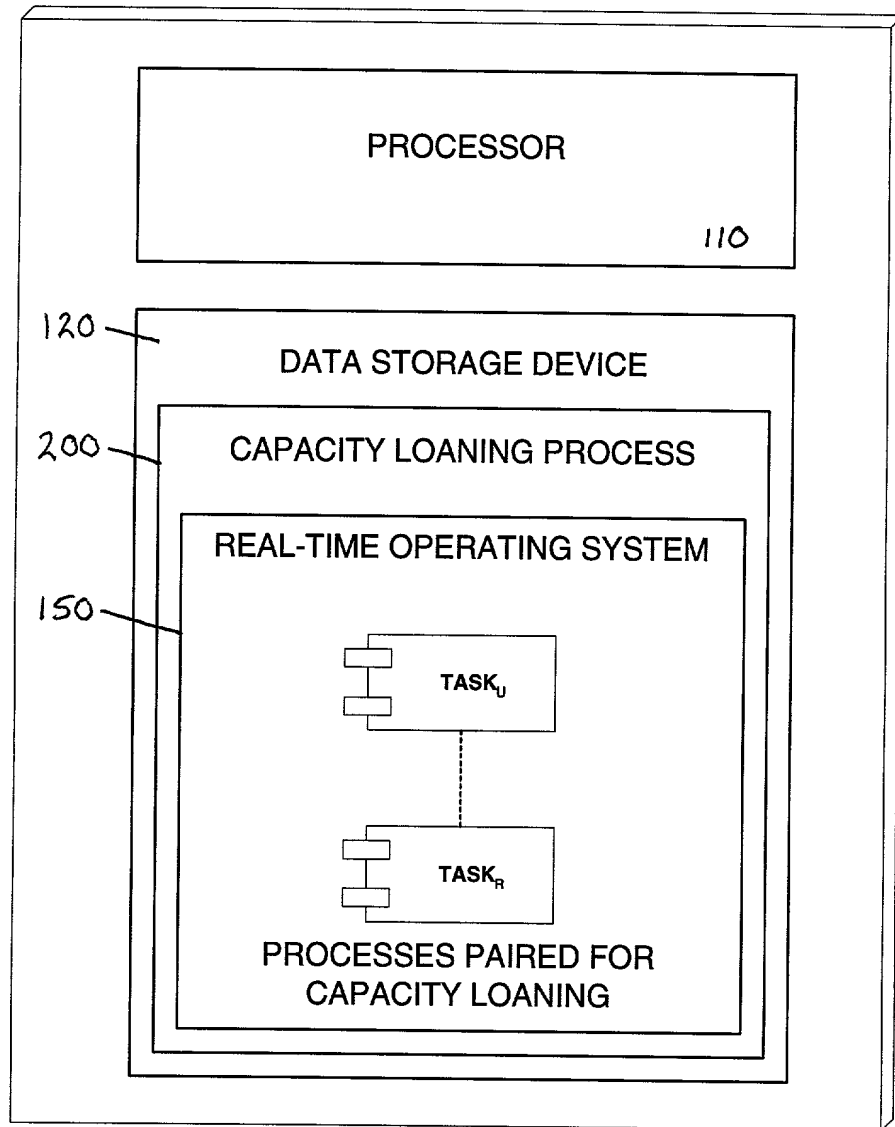


FIG. 1

CAPACITY LOANING PROCESS - 200

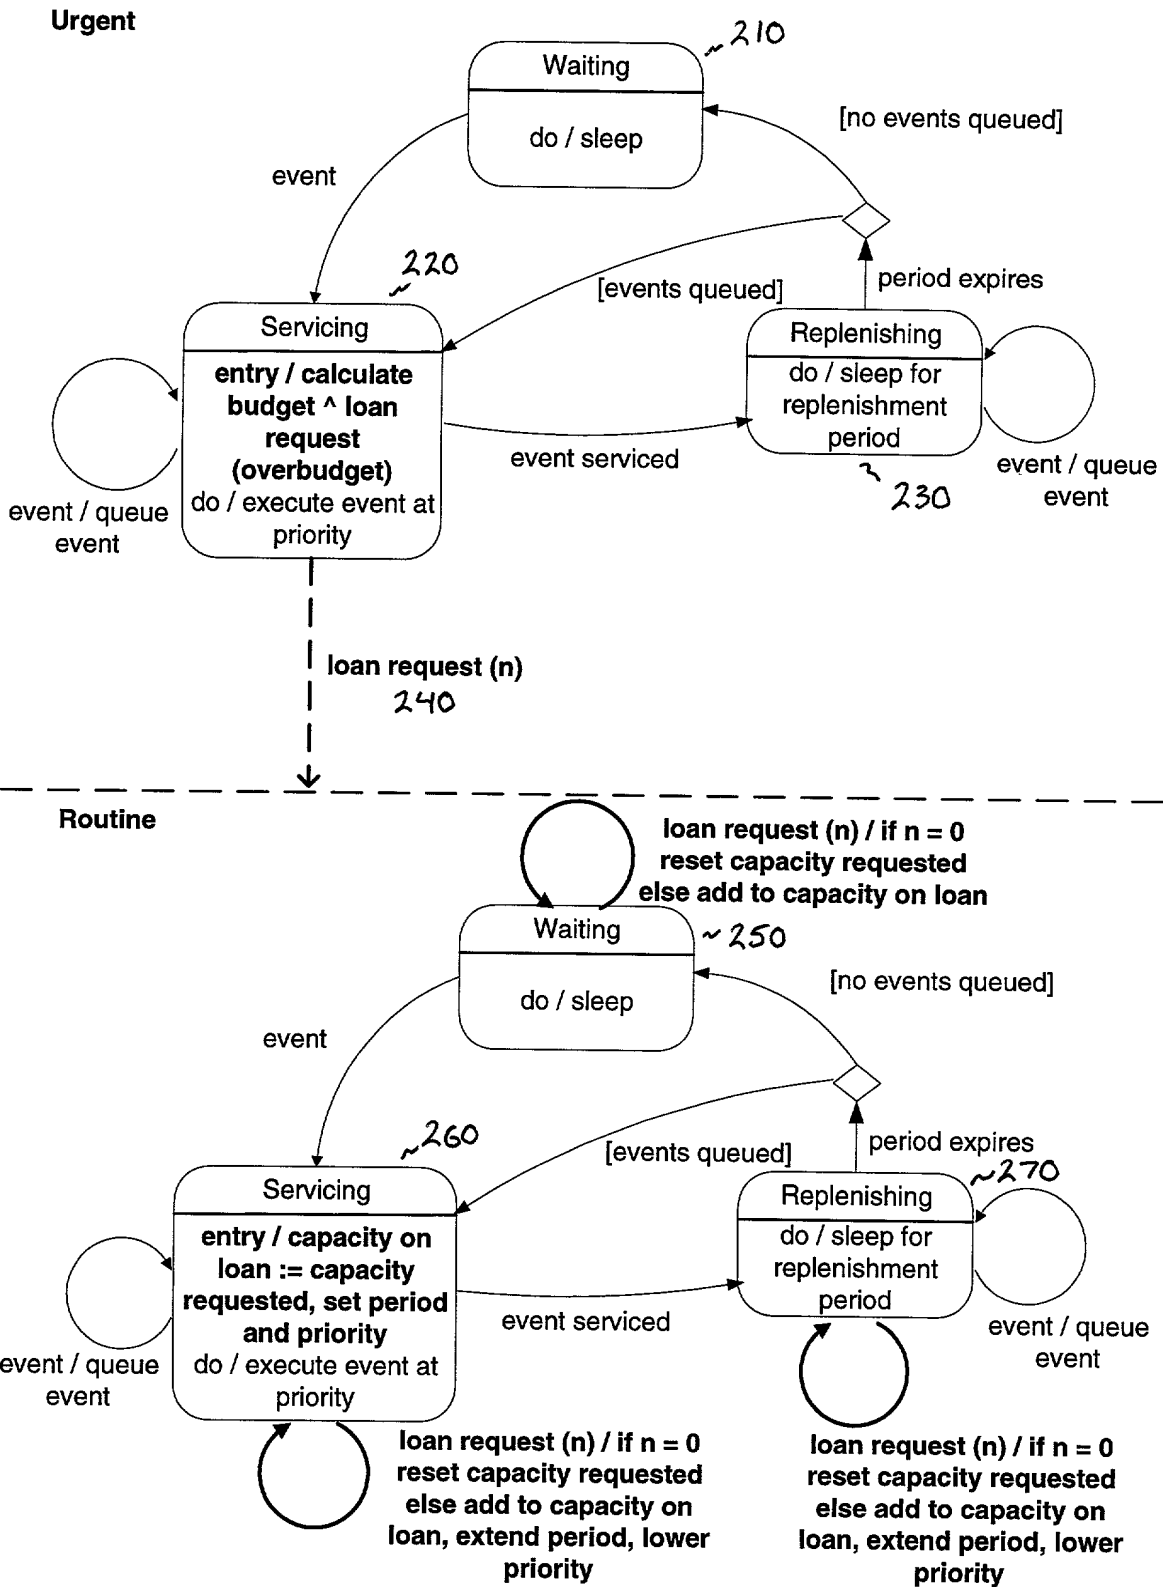


FIG. 2

DECLARATION and POWER OF ATTORNEY

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled **Method for Dynamic Loaning in Rate Monotonic Real-Time Systems** the specification of which (check one)

☒ is attached hereto.

_____ was filed on _____ as Application Serial No. _____ and was amended on _____

(if applicable).

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by the amendment(s) referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulation, §1.56(a).

I hereby claim foreign priority benefits under Title 35, United States Code, § 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

PRIOR FOREIGN APPLICATION(S)

COUNTRY	APPLICATION NUMBER	DATE OF FILING (DAY, MONTH, YEAR)	PRIORITY CLAIMED UNDER 35 U.S.C. 119

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application (s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35 United States Code, § 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, §1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

PRIOR UNITED STATES APPLICATION(S)

APPLICATION SERIAL NUMBER	FILING DATE	STATUS (PATENTED, PENDING, ABANDONED)
60/129,301	April 14, 1999	Pending

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.


POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. (list name and registration number)

Algy Tamoshunas, Reg. No. 27,677

Jack E. Haken, Reg. No. 26,902

SEND CORRESPONDENCE TO: Corporate Patent Counsel; U.S. Philips Corporation; 580 White Plains Road;	DIRECT TELEPHONE CALLS TO: Gregory L. Thorne (914) 333-9665
--	---

Tarrytown, NY 10591

Dated: 20 Dec 1999		Inventor's Signature: 		
Full Name of Inventor	Last Name: Isham	First Name : Karl	Middle Name: Michael	
Residence & Citizenship	City Grantsville	State or Foreign Country Utah	Country of Citizenship United States of America	
Post Office Address	Street P.O. Box 1384	City Grantsville	State or Country Utah	Zip Code 84029

Dated:		Inventor's Signature:		
Full Name of Inventor	Last Name:	First Name :	Middle Name:	
Residence & Citizenship	City	State or Foreign Country	Country of Citizenship	
Post Office Address	Street	City	State or Country	Zip Code

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of

Atty. Docket

KARL M. ISHAM

PHA 23,656

Serial No.

Filed: CONCURRENTLY

Title: METHOD FOR DYNAMIC LOANING IN RATE MONOTONIC REAL-TIME SYSTEMS

Honorable Commissioner of Patents and Trademarks
Washington, D.C. 20231

APPOINTMENT OF ASSOCIATES

Sir:

The undersigned Attorney of Record hereby revokes all prior appointments (if any) of Associate Attorney(s) or Agent(s) in the above-captioned case and appoints:


Tony Piotrowski

(Registration No. 42,080)

c/o U.S. PHILIPS CORPORATION, Intellectual Property Department, 580 White Plains Road, Tarrytown, New York 10591, his Associate Attorney(s)/Agent(s) with all the usual powers to prosecute the above-identified application and any division or continuation thereof, to make alterations and amendments therein, and to transact all business in the Patent and Trademark Office connected therewith.

ALL CORRESPONDENCE CONCERNING THIS APPLICATION AND THE LETTERS PATENT WHEN GRANTED SHOULD BE ADDRESSED TO THE UNDERSIGNED ATTORNEY OF RECORD.

Respectfully,



Jack E. Haken, Reg. 26,902
Attorney of Record

Dated at Tarrytown, New York
this 21st day of December, 1999.